arXiv:1511.00511v1 [cs.PL] 2 Nov 2015

A Formal Model for Direct-style Asynchronous Observables

Philipp Haller[1] and Heather Miller[2]

[1] KTH Royal Institute of Technology, Sweden
phaller@kth.se
[2] EPFL, Switzerland
heather.miller@epfl.ch

Abstract

Languages like F#, C#, and recently also Scala, provide "Async" programming models which aim to make asynchronous programming easier by avoiding an inversion of control that is inherent in callback-based programming models. This paper presents a novel approach to integrate the Async model with observable streams of the Reactive Extensions model. Reactive Extensions are best-known from the .NET platform, and widely-used implementations of its programming model exist also for Java, Ruby, and other languages. This paper contributes a formalization of the unified "Reactive Async" model in the context of an object-based core calculus. Our formal model captures the essence of the protocol of asynchronous observables using a heap evolution property. We prove a subject reduction theorem; the theorem implies that reduction preserves the heap evolution property. Thus, for well-typed programs our calculus ensures the protocol of asynchronous observables.

1 Introduction

Asynchronous programming has been a challenge for a long time. A multitude of programming models have been proposed that aim to simplify the task. Interestingly, there are elements of a convergence arising, at least with respect to the basic building blocks: futures and promises have begun to play a more and more important role in a number of languages like Java, C++, ECMAScript, and Scala.

The Async extensions of F# [7], C# [1], and Scala [4] provide language support for programming with futures (or "tasks"), by avoiding an inversion of control that is inherent in designs based on callbacks. However, these extensions are so far only applicable to futures or future-like abstractions. In this paper we present an integration of the Async model with a richer underlying abstraction, the observable streams of the Reactive Extensions model. [6] A reactive stream is a stream of *observable events* which an arbitrary number of *observers* can subscribe to. The set of possible event patterns of observable streams is strictly greater than those of futures. A stream can (a) produce zero or more regular events, (b) complete normally, or (c) complete with an error (it's even possible for a stream to never complete.) Given the richer substrate of reactive streams, the Async model has to be generalized in several dimensions.

This paper makes the following contributions:

- A design of a new programming model, RAY, which integrates the Async model and the Reactive Extensions model (the name RAY is inspired by the main constructs, `rasync`, `await` and `yield`);
- A formal model of the proposed programming model. Our operational semantics extends the formal model presented in [1] for C#'s async/await to observable streams. The formal model captures the essence of the protocol of asynchronous observables using a heap evolution property;

- A proof of subject reduction for the presented core calculus. The theorem implies that reduction preserves the heap evolution property. Thus, for well-typed programs our calculus ensures the protocol of asynchronous observables.

The rest of the paper is organized as follows. The following Section 2 provides background on Scala Async, an implementation of the Async model, and the Reactive Extensions model. Section 3 introduces our unified Reactive Async model. In Section 4 we present a formalization of the Reactive Async model in the context of an object-based core calculus. Section 5 presents correctness properties including subject reduction.

2 Background

2.1 Scala Async

Scala Async provides constructs that aim to facilitate programming with asynchronous events in Scala. The introduced constructs are inspired to a large extent by extensions that have been introduced in C# version 5 [5] in a similar form. The goal is to enable expressing asynchronous code in "direct style", *i.e.,* in a familiar blocking style where suspending operations look as if they were blocking while at the same time using efficient non-blocking APIs under the hood.

In Scala, an immediate consequence is that non-blocking code using Scala's futures API [3] does not have to resort to (a) low-level callbacks, or (b) higher-order functions like `map` and `flatMap`. While the latter have great composability properties, they can appear unnatural when used to express the regular control flow of a program.

For example, an efficient non-blocking composition of asynchronous web service calls using futures can be expressed as follows in Scala:

```
1  val futureDOY: Future[Response] =
2    WS.url("http://api.day-of-year/today").get
3
4  val futureDaysLeft: Future[Response] =
5    WS.url("http://api.days-left/today").get
6
7  futureDOY.flatMap { doyResponse =>
8    val dayOfYear = doyResponse.body
9    futureDaysLeft.map { daysLeftResponse =>
10     val daysLeft = daysLeftResponse.body
11     Ok("" + dayOfYear + ": " +
12             daysLeft + " days left!")
13   }
14 }
```

Line 1 and 4 define two futures obtained as results of asynchronous requests to two hypothetical web services using an API inspired by the Play! Framework (for the purpose of this example, the definition of type `Response` is unimportant).

This can be expressed more intuitively in direct style using Scala Async as follows (this example is adopted from the SIP proposal [4]):

```
1  val respFut = async {
2    val dayOfYear = await(futureDOY).body
3    val daysLeft = await(futureDaysLeft).body
```

```
4    Ok("" + dayOfYear + ": " +
5            daysLeft + " days left!")
6  }
```

The `await` on line 2 causes the execution of the `async` block to suspend until `futureDOY` is completed (with a successful result or with an exception). When the future is completed successfully, its result is bound to the `dayOfYear` local variable, and the execution of the `async` block is resumed. When the future is completed with an exception (for example, because of a timeout), the invocation of `await` re-throws the exception that the future was completed with. In turn, this completes future `respFut` with the same exception. Likewise, the `await` on line 3 suspends the execution of the `async` block until `futureDaysLeft` is completed.

The main methods provided by Scala Async, `async` and `await`, have the following type signatures:

```
def async[T](body: => T): Future[T]
def await[T](future: Future[T]): T
```

Given the above definitions, `async` and `await` "cancel each other out:"

```
await(async { <expr> }) = <expr>
```

This "equation" paints a grossly over-simplified picture, though, since the actual operational behavior is much more complicated: `async` typically schedules its argument expression to run asynchronously on a thread pool; moreover, `await` may only be invoked within a syntactically enclosing `async` block.

2.2 Reactive Extensions

The Rx programming model is based on two interface traits: `Observable` and `Observer`. `Observable` represents observable streams, *i.e.,* streams that produce a sequence of events. These events can be observed by registering an `Observer` with the `Observable`. The `Observer` provides methods which are invoked for each of the kinds of events produced by the `Observable`. In Scala, the two traits can be defined as shown in Figure 1.

```
trait Observable[T] {
  def subscribe(obs: Observer[T]): Closable
}

trait Observer[T] extends (Try[T] => Unit) {
  def apply(tr: Try[T]): Unit
  def onNext(v: T) = apply(Success(v))
  def onFailure(t: Throwable) = apply(Failure(t))
  def onDone(): Unit
}
```

Figure 1: The `Observable` and `Observer` traits.

The idea of the `Observer` is that it can respond to three different kinds of events, (1) the next regular event (`onNext`), (2) a failure (`onFailure`), and (3) the end of the observable stream (`onDone`). Thus, the two traits constitute a variation of the classic subject/observer pattern [2]. Note that `Observable`'s `subscribe` method returns a `Closable`; it has only a single abstract

`close` method which removes the subscription from the observable. The next listing shows an example implementation.

Note that in our Scala version the `Observer` trait extends the function type `Try[T] => Unit`. `Try[T]` is a simple container type which supports heap-based exception handling (as opposed to the traditional stack-based exception handling using expressions like `try-catch-finally`.) There are two subclasses of `Try[T]`: `Success` (encapsulating a value of type `T`) and `Failure` (encapsulating an exception). Given the above definition, a concrete `Observer` only has to provide implementations for the `apply` and `onDone` methods. Since `apply` takes a parameter of type `Try[T]` its implementation handles the `onNext` and `onFailure` events all at once (in Scala, this is tyically done by pattern matching on `tr` with cases for `Success` and `Failure`).

The `Observer` and `Observable` traits are used as follows. For example, here is a factory method for creating an observable from a text input field of typical GUI toolkits (this example is adapted from [6]):

```
def textChanges(tf: JTextField): Observable[String] =
  new ObservableBase[String] {
    def subscribe(o: Observer[String]) = {
      val l = new DocumentListener {
        def changedUpdate(e: DocumentEvent) = {
          o.onNext(tf.getText())
        }
      }
      tf.addDocumentListener(l)
      new Closable() {
        def close() = {
          tf.removeDocumentListener(l)
        }
      }
    }
  }
```

This newly-defined `textChanges` combinator can be used with other Rx combinators as follows:

```
textChanges(input)
.flatMap(word => completions(word))
.subscribe(observeChanges(output))
```

We start with the observable created using the `textChanges` method from above. Then we use the `flatMap` combinator (called `Select` in C#) to transform the observable into a new observable which is a stream of completions for a given word (a string). On the resulting observable we call `subscribe` to register a consumer: `observeChanges` creates an observer which outputs all received events to the `output` stream. (The shown example suffers from a problem explained in [6] which motivates the use of an additional `Switch` combinator which is omitted here for brevity.)

3 The Reactive Async Model

This Section provides an (example-driven) overview of the Reactive Async Model which integrates the Async Model and the Reactive Extensions Model.

The basic idea is to generalize the Async model, so that it can be used not only with futures, but also with observable streams. This means, we need constructs that can create observables, as opposed to only futures (like `async`), and we need ways to wait for more events than just the completion of a future. Essentially, it should be possible to await all kinds of events produced by an observable stream. Analogous to `await` which waits for the completion event of a future, we introduce variations like `awaitNext` and `awaitNextOrDone` to express waiting for the events of an observable stream.

3.1 Example

The following example shows how to await a fixed number of events of a stream in the Reactive Async Model:

```
val obs = rasync {
  var events = List[Int]()
  while (events.size < 5) {
    val event = awaitNext(stream)
    events = event :: events
  }
  Some(events)
}
```

Note that we are using the `rasync` construct; it is a generalized version of the `async` construct of Section 2.1 which additionally supports methods to await events of observable streams.

In the above example, the invocation of `awaitNext` suspends the `rasync` block until the producer of `stream` calls `onNext` on its observers. The argument of this `onNext` call (the next event) is returned as a result from `awaitNext`. The result of `rasync`, `obs`, has type `Observable[List[Int]]`. Once the body of `rasync` has been fully evaluated, `obs` publishes two events: first, an `onNext` event which carries `events` (the list with five elements), and second, an `onDone` event; it is not possible for `obs` to publish further events.

Note that the result of an `rasync` block has a type of the form `Option[T]`; in the case where this optional value is empty (`None`), only an `onDone` event is published as a result of fully evaluating the `rasync` block. (It is, however, possible to publish other events beforehand, as shown in the following sections.) Otherwise, the semantics of `rasync` is analogous to the behavior of a regular `async` block: when its body has been fully evaluated, the future, which is the result of `async`, is completed and further changes to the state of the future are impossible.

3.2 Awaiting Stream Termination

Sometimes it is not known statically how many events a stream might still publish. One might want to collect all events until the stream is done (finished publishing events). In this case it is necessary to have a way to wait for either of two events: the stream publishes a next event, or the stream is done. This can be supported using a method `awaitNextOrDone` which returns an `Option[T]` when applied to an `Observable[T]`:

```
rasync {
  var events: List[Int] = List()
  var next: Option[Int] = awaitNextOrDone(stream)
  while (next.nonEmpty) {
    events = next.get :: events
```

```
    next = awaitNextOrDone(stream)
  }
  Some(events)
}
```

In the above example, the body of `rasync` repeatedly waits for the given `stream` to publish either a next event or to reach its end, using `awaitNextOrDone`. As long as the `stream` continues to publish events (in which case `next` of type `Option[Int]` is non-empty), each event is prepended to the `events` list; this list is the single event that the observable which is, in turn, created by `rasync` publishes (once the body of `rasync` has been fully evaluated).

3.3 Creating Complex Streams

The streams created by `rasync` in the previous sections are rather simple: after consuming events from other streams only a single interesting event is published on the created stream (by virtue of reaching the end of the `rasync` block). In this section, we explain how more complex streams can be created in the Reactive Async Model.

```
val forwarder = rasync[Int] {
  var next: Option[Int] =
    awaitNextOrDone(stream)

  while (next.nonEmpty) {
    yieldNext(next.get)
    next = awaitNextOrDone(stream)
  }
  None
}
```

Figure 2: A simple forwarder stream.

Suppose we would like to create a stream which simply publishes an event for each event observed on another `stream`. In this case, the constructs we have seen so far are not sufficient, since an arbitrary number of events have to be published from within the `rasync` block. This is where the new method `yieldNext` comes in: it publishes the next event to the stream returned by `rasync`. Our simple forwarder example can then be expressed as shown in Figure 2.

Note that in the above example, the result of the body of the `rasync` block is `None`; consequently, the resulting `forwarder` stream only publishes an `onDone` event when `rasync`'s body has been fully evaluated. In this case, it is assumed that the only "interesting" non-done events of `forwarder` are published using `yieldNext`.

4 Formalization

One of the contributions of this paper is an operational semantics of the proposed programming model. Our operational semantics generalizes the formal model presented in [1].

$p ::= \overline{cd}\ e$	program
$cd ::= \texttt{class}\ C\ \{\overline{fd}\ \overline{md}\}$	class declaration
$fd ::= \texttt{var}\ f : \sigma$	field declaration
$md ::= \texttt{def}\ m(\overline{x : \sigma}) : \tau = e$	method declaration
$\sigma, \tau ::=$	type
$\quad \mid \gamma$	value type
$\quad \mid \rho$	reference type
$\gamma ::=$	value type
$\quad \mid \texttt{Boolean}$	boolean
$\quad \mid \texttt{Int}$	integer
$\rho ::=$	reference type
$\quad \mid C$	class type
$\quad \mid \texttt{Observable}[\sigma]$	observable type

Figure 3: Core language syntax. C is a class name, f, m are field and method names.

$e ::=$	expressions
$\quad \mid \underline{b}$	boolean
$\quad \mid \underline{i}$	integer
$\quad \mid x$	variable
$\quad \mid \texttt{null}$	null
$\quad \mid \texttt{if}\ (x)\ \{e\}\ \texttt{else}\ \{e'\}$	condition
$\quad \mid \texttt{while}\ (x)\ \{e\}$	while loop
$\quad \mid x.f$	selection
$\quad \mid x.f = y$	assignment
$\quad \mid x.m(\overline{y})$	invocation
$\quad \mid \texttt{new C}(\overline{y})$	instance creation
$\quad \mid \texttt{let}\ x = e\ \texttt{in}\ e'$	let binding
$\quad \mid \texttt{rasync}[\sigma](\bar{y})\ \{e\}$	observable creation
$\quad \mid \texttt{await}(x)$	await next event
$\quad \mid \texttt{yield}(x)$	yield event

Figure 4: RAY expressions.

4.1 Syntax

Figure 3 and Figure 4 show the syntax of our core language. Note that programs are written in *A-normal form* (ANF) which forces all subexpressions to be named; this simplifies the presentation of the operational semantics. Note that our core language does not support any form of subtyping; thus, class declarations do not specify a superclass. This is adopted from [1]; the presented reactive features are orthogonal to subtyping.

A RAY program consists of a collection of class definitions, as well as the body (an expression) of a "main" method. A class `C` has (a possibly empty) sequence of public fields and methods, $\overline{f}$ with types $\overline{\sigma}$, and $\overline{md}$, respectively. Methods are public with return type τ, a type which may represent either a value or reference type, and their body is an expression e.

The `Observable[`σ`]` family of types is used to model the generic nature of observables. They represent observables such that an `rasync` block can await their next event using `await`. Conversely, inside the body e of an expression `rasync[`σ`]`$(\bar{y})\ \{e\}$, `yield` can be used to publish events of type σ.

4.2 Operational Semantics

4.2.1 Notation

A heap, denoted H, partially maps object identifiers (ranged over by o) to heap objects, denoted $\langle \texttt{C}, FM \rangle$, representing a pair of type `C` and a field map, FM. A field map partially maps fields f to values (ranged over by v), where v can be either an integer, a boolean, `null`, or an object identifier (the address of an object in the heap).

Frames have the form $\langle L, e \rangle^l$ where L maps local variables to their values, e is an expression, and l is a label. A label is either s denoting a regular, synchronous frame, or $a(o, \bar{p})$ denoting an asynchronous frame; in this case, o is the heap address of a corresponding observable object $\langle \texttt{Observable}[\sigma], running(\bar{F}, \bar{S}) \rangle$; $\bar{p}$ is a sequence of object identifiers of observables that observable o has subscribed to. $\bar{F}$ is a set of asynchronous frames, namely, all observables that are currently suspended awaiting o to publish a new event. $\bar{S}$ is a set of subscribers which are explained below.

There are three kinds of transition rules. The first kind goes from a heap and a frame to a new heap and a new frame (simple right arrow). The second kind goes from a heap and a frame stack to a new heap and a new frame stack (double right arrow). The third kind goes from a heap and a set of frame stacks to a new heap and a new set of frame stacks (squiggly right arrow).

4.2.2 Synchronous Transition Rules

$$\frac{}{H, \langle L, \texttt{let } x = y \texttt{ in } e \rangle^l \longrightarrow H, \langle L[x \mapsto L(y)], e \rangle^l} \quad \text{(E-Var)}$$

$$\frac{H(L(y)) = \langle \rho, FM \rangle}{H, \langle L, \texttt{let } x = y.f \texttt{ in } e \rangle^l \longrightarrow H, \langle L[x \mapsto FM(f)], e \rangle^l} \quad \text{(E-Field)}$$

$$\frac{F' = \begin{cases} \langle L, \texttt{let } x' = e \texttt{ in let } x = \texttt{while } (y) \ \{e\} \texttt{ in } e' \rangle^l & \text{if } L(y) = \texttt{true} \qquad x' \notin dom(L) \\ \langle L, \texttt{let } x = \texttt{false in } e' \rangle^l & \text{if } L(y) = \texttt{false} \end{cases}}{H, \langle L, \texttt{let } x = \texttt{while } (y) \ \{e\} \texttt{ in } e' \rangle^l \longrightarrow H, F'} \quad \text{(E-While)}$$

$$\frac{F' = \begin{cases} \langle L, \texttt{let } x = t \texttt{ in } u \rangle^l & \text{if } L(y) = \texttt{true} \\ \langle L, \texttt{let } x = s \texttt{ in } u \rangle^l & \text{if } L(y) = \texttt{false} \end{cases}}{H, \langle L, \texttt{let } x = \texttt{if } (y) \ \{t\} \texttt{ else } \{s\} \texttt{ in } u \rangle^l \longrightarrow H, F'} \quad \text{(E-Cond)}$$

$$\frac{L(x) = o \qquad H(o) = \langle \sigma, FM \rangle \qquad H' = H[o \mapsto \langle \sigma, FM[f \mapsto L(y)] \rangle]}{H, \langle L, \texttt{let } x' = x.f = y \texttt{ in } e \rangle^l \longrightarrow H', \langle L, \texttt{let } x' = y \texttt{ in } e \rangle^l} \quad \text{(E-Assign)}$$

$$\frac{\begin{array}{c} fields(C) = \bar{f} \qquad o \notin dom(H) \\ H' = H[o \mapsto \langle C, \bar{f} \mapsto L(\bar{y}) \rangle] \end{array}}{H, \langle L, \texttt{let } x = \texttt{ new } C(\bar{y}) \texttt{ in } e \rangle^l \longrightarrow H', \langle L[x \mapsto o], e \rangle^l} \quad \text{(E-New)}$$

Figure 5: Simple frame transition rules.

Figure 5 shows simple frame transition rules. Note that all transition rules preserve the

labels of frames. Rule E-Var looks up the value $L(y)$ of y in local variable mapping L; the local variable mapping of the target frame maps x to $L(y)$; reduction continues with expression e. Rule E-Field looks up the value of field $y.f$ using L and H; as before, reduction continues with e. Rule E-While looks up the value of the condition variable in the local variable mapping and continues reduction accordingly. Note that our core language does not distinguish between expressions and statements; thus, also the while loop is an expression which reduces to the constant `false` when the loop condition is `false`. Rule E-Cond is standard. Rule E-Assign combines the heap H, local variable mapping L, and field mapping FM in the natural way for field assignment. Rule E-New creates a new instance of class C, assigning the constructor arguments $\bar{y}$ to the fields of the new instance.

$$\frac{H(L(y)) = \langle \rho, FM \rangle \qquad mbody(\rho, m) = (\overline{x}) \to e' \qquad L' = [\bar{x} \mapsto L(\bar{z}), \texttt{this} \mapsto L(y)]}{H, \langle L, \texttt{let } x = y.m(\bar{z}) \texttt{ in } e \rangle^l \circ FS \;\twoheadrightarrow\; H, \langle L', e' \rangle^s \circ \langle L, e \rangle^l{}_x \circ FS} \quad \text{(E-Method)}$$

$$\frac{}{H, \langle L, y \rangle^s \circ \langle L', e \rangle^l{}_x \circ FS \;\twoheadrightarrow\; H, \langle L'[x \mapsto L(y)], e \rangle^l \circ FS} \quad \text{(E-Return)}$$

$$\frac{H, F \;\longrightarrow\; H', F'}{H, F \circ FS \;\twoheadrightarrow\; H', F' \circ FS} \quad \text{(E-Frame)}$$

Figure 6: Method call and return transition rules.

Figure 6 shows the transition rules for method call and return. Rule E-Method evaluates a method invocation. The run-time type of the receiver, ρ, is looked up in heap H. Using the auxiliary function *mbody* we look up the body of method m in ρ. To evaluate the method body, a new frame with synchronous label s is created and pushed on top of the frame stack. Importantly, the caller frame (with expression e) is annotated with variable x; this annotation is used for the transfer of the return value as follows. Rule E-Return shows how a value is returned from a method invocation to the caller. A method call returns when the expression of its frame has been reduced to a variable y. The method's frame is popped off the frame stack, and the frame of the caller is replaced with a frame that maps variable x to the value of y. Crucially, the frame of the caller is annotated with x. Finally, rule E-Frame transitions a *frame stack* $F \circ FS$ by transitioning frame F.

4.2.3 Asynchronous Transition Rules

Figure 7 shows the asynchronous transition rules. These rules transition either between frame stacks ($\twoheadrightarrow$) or processes ($\rightsquigarrow$).

Rule E-RAsync evaluates the creation of a new observable using an expression of the form `rasync[`σ`]`$(\bar{y})$ $\{e\}$. The type argument σ determines the type of the observable, `Observable`$[\sigma]$. The value arguments $\bar{y}$ refer to observables at addresses $\bar{p}$ in heap H.

Recall the representation of observables in the heap. A newly-created observable has the form $\langle \texttt{Observable}[\sigma], running(\epsilon, \epsilon) \rangle$. In general, a running observable (*i.e.,* an observable that has not terminated, yet) has the form $\langle \psi, running(\bar{F}, \bar{S}) \rangle$. Each *waiter* $F \in \bar{F}$ is an asynchronous frame of an observable waiting to receive an event. Each *subscriber* $S \in \bar{S}$ is a pair $\langle o, q \rangle$; o is the address of a subscribed observable that is running, but currently not waiting to receive an event; q is a queue of events received by o (but not yet processed).

Rule E-RAsync adds the newly-created observable o as a subscriber (with an empty queue)

$$\frac{\begin{array}{c}L(\bar{y}) = \bar{p} \qquad o \notin dom(H) \\ H_0 = H[o \mapsto \langle \texttt{Observable}[\sigma], running(\epsilon, \epsilon)\rangle] \\ \forall p_i \in \bar{p}.\ H(p_i) = \langle \psi_i, running(\bar{F}_i, \bar{S}_i)\rangle \\ \forall i \in 1 \ldots n.\ H_i = H_{i-1}[p_i \mapsto \langle \psi_i, running(\bar{F}_i, \langle o, []\rangle :: \bar{S}_i)\rangle] \\ H' = H_n\end{array}}{\begin{array}{c}H, \langle L, \texttt{let } x = \texttt{rasync}[\sigma]\,(\bar{y})\ \{e\} \texttt{ in } t\rangle^l \circ FS \\ \twoheadrightarrow\ H', \langle L, e\rangle^{a(o,\bar{p})} \circ \langle L[x \mapsto o], t\rangle^l \circ FS\end{array}} \text{(E-RAsync)}$$

$$\frac{\begin{array}{c}H(o) = \langle \texttt{Observable}[\sigma], running(\bar{F}, \bar{S})\rangle \\ \bar{R} = resume(\bar{F}, \texttt{Some}(L(z))) \qquad Q = \{R \circ \epsilon \mid R \in \bar{R}\} \\ \bar{S}' = \{\langle o', L(z) :: q\rangle \mid \langle o', q\rangle \in \bar{S}\} \\ H' = H[o \mapsto \langle \texttt{Observable}[\sigma], running(\epsilon, \bar{S}')\rangle]\end{array}}{H, \{\langle L, \texttt{yield}(z)\rangle^{a(o,\bar{p})} \circ FS\} \cup P\ \rightsquigarrow\ H', \{\langle L, z\rangle^{a(o,\bar{p})} \circ FS\} \cup P \cup Q} \text{(E-Yield)}$$

$$\frac{\begin{array}{c}H(o) = \langle \texttt{Observable}[\sigma], running(\bar{F}, \bar{S})\rangle \\ \bar{R} = resume(\bar{F}, \texttt{None}) \qquad Q = \{R \circ \epsilon \mid R \in \bar{R}\} \\ H_0 = H[o \mapsto \langle \texttt{Observable}[\sigma], done(\bar{S})\rangle] \\ \forall i \in 1 \ldots n.\ H_i = H_{i-1}[p_i \mapsto unsub(o, p_i, H)]\end{array}}{H, \{\langle L, x\rangle^{a(o,\bar{p})} \circ FS\} \cup P\ \rightsquigarrow\ H_n, \{FS\} \cup P \cup Q} \text{(E-RAsync-Return)}$$

$$\frac{\begin{array}{c}F = \langle L, \texttt{let } x = \texttt{ await}(y) \texttt{ in } t\rangle^{a(o,\bar{p})} \qquad L(y) = o' \\ H(o') = \langle \texttt{Observable}[\sigma], running(\bar{F}, \bar{S})\rangle \qquad \bar{S} = \bar{R} \uplus \{\langle o, []\rangle\} \\ H' = H[o' \mapsto \langle \texttt{Observable}[\sigma], running(F :: \bar{F}, \bar{R})\rangle]\end{array}}{H, F \circ FS\ \twoheadrightarrow\ H', FS} \text{(E-Await1)}$$

$$\frac{\begin{array}{c}L(y) = o' \qquad H(o') = \langle \texttt{Observable}[\sigma], running(\bar{F}, \bar{S})\rangle \qquad \bar{S} = \bar{R} \uplus \{\langle o, q :: v\rangle\} \\ H' = H[o' \mapsto \langle \texttt{Observable}[\sigma], running(\bar{F}, \bar{R} \cup \{\langle o, q\rangle\})\rangle]\end{array}}{H, \langle L, \texttt{let } x = \texttt{ await}(y) \texttt{ in } t\rangle^{a(o,\bar{p})} \circ FS\ \twoheadrightarrow\ H', \langle L[x \mapsto \texttt{Some}(v)], t\rangle^{a(o,\bar{p})} \circ FS} \text{(E-Await2)}$$

$$\frac{\begin{array}{c}L(y) = o' \qquad H(o') = \langle \texttt{Observable}[\sigma], done(\bar{S})\rangle \qquad \bar{S} = \bar{R} \uplus \{\langle o, q :: v\rangle\} \\ H' = H[o' \mapsto \langle \texttt{Observable}[\sigma], done(\bar{R} \cup \{\langle o, q\rangle\})\rangle]\end{array}}{H, \langle L, \texttt{let } x = \texttt{ await}(y) \texttt{ in } t\rangle^{a(o,\bar{p})} \circ FS\ \twoheadrightarrow\ H', \langle L[x \mapsto \texttt{Some}(v)], t\rangle^{a(o,\bar{p})} \circ FS} \text{(E-Await3)}$$

Figure 7: Asynchronous transition rules.

$$resume(\bar{F}, v) = \{\langle L[x \mapsto v], t\rangle^{a(o,\bar{p})} \mid \langle L, \texttt{let } x = \texttt{await}(y) \texttt{ in } t\rangle^{a(o,\bar{p})} \in \bar{F}\}$$

$$unsub(\bar{S}, o) = \{\langle o', q\rangle \mid \langle o', q\rangle \in \bar{S} \wedge o' \neq o\}$$

Figure 8: Auxiliary functions.

to each observable $p_i \in \bar{p}$. Finally, a new asynchronous frame $\langle L, e\rangle^{a(o,\bar{p})}$ is pushed on to the frame stack.

Rule E-Yield implements the built-in `yield` expression which publishes a new event to

all waiters and subscribers of observable o. Waiters $F \in \bar{F}$ are resumed by creating a set of new frame stacks Q based on frames created using the *resume* auxiliary function (defined in Figure 8). Rule E-YIELD resumes waiters with value $L(z)$. The state of observable o is updated such that the set of waiters is empty. The `yield` expression itself is reduced to z.

Rule E-RASYNC-RETURN transitions a process where an asynchronous frame $\langle L, x\rangle^{a(o,\bar{p})}$ is on top of one of the frame stacks. Since the frame's expression has been reduced to just a variable x, this means the corresponding observable o is done publishing events at this point. Therefore, the waiters $\bar{F}$ are resumed with result `None`. The state of observable o is updated to $\langle \texttt{Observable}[\sigma], \mathit{done}(\bar{S})\rangle$; value $\mathit{done}(\bar{S})$ indicates that o has transitioned to the terminated state; events published to subscribers ($\bar{S}$) remain available for consumption, though (see rules for `await` below). Finally, o unsubscribes from all observables $\bar{p}$. $\mathit{unsub}(o, p, H)$ is defined as follows; the function makes use of $\mathit{unsub}(\bar{S}, o)$ which is defined in Figure 8.

$$\mathit{unsub}(o,p,H) = \begin{cases} \langle \psi, \mathit{running}(\bar{F}, \mathit{unsub}(\bar{S}, o))\rangle & \text{if } H(p) = \langle \psi, \mathit{running}(\bar{F}, \bar{S})\rangle \\ \langle \psi, \mathit{done}(\mathit{unsub}(\bar{S}, o))\rangle & \text{if } H(p) = \langle \psi, \mathit{done}(\bar{S})\rangle \end{cases}$$

Rules E-AWAIT1, E-AWAIT2, and E-AWAIT3 implement the built-in `await` expression. Rule E-AWAIT1 adds the asynchronous frame F of observable o to the waiters of observable o' in the case where there is no event from o' ready to be consumed by o. Rule E-AWAIT2 handles the dual case where observable o immediately receives an event from observable o'; the subscribers of o' are updated accordingly in the target heap H'. Rule E-AWAIT3 handles the case where observable o' is in a terminated state $\mathit{done}(\bar{S})$. Importantly, a subscriber queue in $\bar{S}$ may contain an event to be consumed by the `await`-invoking observable o. In case the corresponding subscriber queue is empty, observable o suspends analogous to rule E-AWAIT1.

$$\frac{}{H, \{\epsilon\} \cup P \rightsquigarrow H, P} \qquad \text{(E-EXIT)}$$

$$\frac{H, FS \twoheadrightarrow H', FS'}{H, \{FS\} \cup P \rightsquigarrow H', \{FS'\} \cup P} \qquad \text{(E-SCHEDULE)}$$

Figure 9: Process transition rules.

Process transition rules enable reducing frame stacks, *i.e.,* threads; Figure 9 shows the transition rules. We use an interleaved semantics. Rule E-SCHEDULE non-deterministically selects and transitions a thread; note that the transition may have side effects on the heap. Rule E-EXIT removes threads with empty frame stacks from the soup of threads.

5 Correctness Properties

We show that well-typed programs satisfy desirable properties:

1. *Observable protocol.* For example, a terminated observable never publishes events again; this protocol property is captured by a *heap evolution* invariant.
2. *Subject reduction.* Reduction of well-typed programs preserves types.

The proofs of these properties are based on a typing relation, as well as invariants preserved by reduction. To establish the correctness properties we have to consider non-interference properties for processes, frame stacks, frames, and heaps; these properties are shown in Figure 10.

The application $obsIds(FS)$ of the $obsIds$ auxiliary function returns the set of all object addresses o in labels $a(o, \bar{p})$ of the asynchronous frames FS (similarly for a single frame F). The $waiters$ function returns the observable ids of the waiting frames of the running state of a given observable heap object. For an observable heap object $H(o) = \langle \texttt{Observable}[\sigma], running(\bar{F}, \bar{S})\rangle$, $waiters(H(o)) = obsIds(\bar{F})$. To test whether an observable is currently running (as opposed to done) we use a simple predicate, $Running$. Finally, to express disjointness of (sets of) heap addresses we use the symbol #.

$$\frac{}{H \vdash \epsilon\ \mathbf{ok}} \quad \text{(EmpFS-ok)}$$

$$\frac{H \vdash F\ \mathbf{ok} \quad H \vdash FS\ \mathbf{ok} \quad obsIds(F) \# obsIds(FS)}{H \vdash F \circ FS\ \mathbf{ok}} \quad \text{(FS-ok)}$$

$$\frac{}{H \vdash F^s\ \mathbf{ok}} \quad \text{(SF-ok)}$$

$$\frac{H \vdash F^s\ \mathbf{ok}}{H \vdash F^s_x\ \mathbf{ok}} \quad \text{(CSF-ok)}$$

$$\frac{\begin{array}{c} Running(H(o)) \\ \forall o' \in dom(H).\ o \notin waiters(H(o')) \wedge (o \in subscribers(H(o')) \Rightarrow o' \in \bar{p}) \end{array}}{H \vdash F^{a(o,\bar{p})}\ \mathbf{ok}} \quad \text{(AF-ok)}$$

$$\frac{\begin{array}{c} \forall o \in dom(H).\ H \vdash H(o)\ \mathbf{ok} \\ \forall o_1 \neq o_2 \in dom(H).\ waiters(H(o_1)) \# waiters(H(o_2)) \end{array}}{\vdash H\ \mathbf{ok}} \quad \text{(H-ok)}$$

$$\frac{}{H \vdash \langle C, FM\rangle\ \mathbf{ok}} \quad \text{(HO-ok)}$$

$$\frac{}{H \vdash \langle \texttt{Observable}[\sigma], done(\bar{S})\rangle\ \mathbf{ok}} \quad \text{(DOHO-ok)}$$

$$\frac{\begin{array}{c} \forall i \neq j \in \{1..n\}.\ obsIds(F_i) \# obsIds(F_j) \\ \forall i \in \{1..n\}.\ \forall o \in obsIds(F_i).\ Running(H(o)) \end{array}}{H \vdash \langle \texttt{Observable}[\sigma], running(F_1, \ldots, F_n, \bar{S})\rangle\ \mathbf{ok}} \quad \text{(ROHO-ok)}$$

$$\frac{\begin{array}{c} H \vdash FS_1\ \mathbf{ok} \quad \ldots \quad H \vdash FS_n\ \mathbf{ok} \\ \forall i \neq j \in \{1..n\}.\ obsIds(FS_i) \# obsIds(FS_j) \end{array}}{H \vdash \{FS_1, \ldots, FS_n\}\ \mathbf{ok}} \quad \text{(Proc-ok)}$$

Figure 10: Non-interference properties.

Rule FS-ok requires that the observable ids in a frame stack are distinct. Rule AF-ok requires that an observable id o in the label of an asynchronous frame is not included in the observable ids of the waiters of any observable in the heap; moreover, if o is included in the observable ids of the subscribers of another observable o', then o' must be included in the set of subscriptions of o, $\bar{p}$. Rule H-ok requires that for all observable objects in the heap, the observable ids of the waiters are disjoint. Rule ROHO-ok requires that a running observable object (a) has no duplicate observable ids in its waiters, and (b) has only waiters that refer to

running (non-terminated) observables. Finally, rule Proc-ok requires that the observable ids in the frame stacks of a process are pairwise disjoint.

To enforce non-interference during evaluation we define a relation between heaps. The following relation also (a) preserves the types of heap objects and (b) bounds the observable ids of new running states.

Definition 1 (Heap Evolution)**.** Heap H evolves to H' wrt a set of observable ids B, written $H \leq_B H'$ if

(i) $\forall o \in dom(H')$. if $o \notin dom(H)$ and $H'(o) = \langle \psi, running(\bar{F}, \bar{S})\rangle$ then $\bar{F} = \bar{S} = \epsilon$, and

(ii) $\forall o \in dom(H)$.

- if $H(o) = \langle C, FM\rangle$ then $H'(o) = \langle C, FM'\rangle$,
- if $H(o) = \langle \psi, done(\bar{S})\rangle$ then $H'(o) = \langle \psi, done(\bar{R} \uplus \{\langle o', q'\rangle\})\rangle$ where $\bar{S} = \bar{R} \uplus \{\langle o', q\rangle\}$, and
- if $H(o) = \langle \psi, running(\bar{F}, \bar{S})\rangle$ then $H'(o) = \langle \psi, running(\bar{F}, \bar{S} \uplus \{\langle o, []\rangle\})\rangle$ or $(H'(o) = \langle \psi, running(\epsilon, \bar{R})\rangle$ and $dom(\bar{S}) = dom(\bar{R}))$ or $(H'(o) = \langle \psi, running(\bar{F} \cup \bar{G}, \bar{S})\rangle$, $obsIds(\bar{F}) \# obsIds(\bar{G})$ and $obsIds(\bar{G}) \subseteq B)$ or $H'(o) = \langle \psi, done(\bar{S})\rangle$.

5.1 Subject Reduction

The following subject reduction theorem is based on a typing relation that is given for processes, frame stacks, frames, expressions, and heaps. Figure 11 shows the typing relation for expressions. It is straight-forward to extend this relation to frames, frame stacks, and processes.

Theorem 1 (Subject Reduction)**.** *If* $\vdash H : \star$ *and* $\vdash H$ ***ok*** *then:*

1. *If* $H \vdash F : \sigma$*,* $H \vdash F$ ***ok*** *and* $H, F \longrightarrow H', F'$ *then* $\vdash H' : \star$*,* $\vdash H'$ ***ok****,* $H' \vdash F' : \sigma$*,* $H' \vdash F'$ ***ok****, and* $\forall B.\ H \leq_B H'$*.*
2. *If* $H \vdash FS : \sigma$*,* $H \vdash FS$ ***ok*** *and* $H, FS \twoheadrightarrow H', FS'$ *then* $\vdash H' : \star$*,* $\vdash H'$ ***ok****,* $H' \vdash FS' : \sigma$*,* $H' \vdash FS'$ ***ok*** *and* $H \leq_{obsIds(FS)} H'$*.*
3. *If* $H \vdash P : \star$*,* $H \vdash P$ ***ok*** *and* $H, P \rightsquigarrow H', P'$ *then* $\vdash H' : \star$*,* $\vdash H'$ ***ok****,* $H' \vdash P' : \star$ *and* $H' \vdash P'$ ***ok****.*

Proof. Part (1) is proved by induction on the derivation of $H, F \longrightarrow H', F'$. Part (2) is proved by induction on the derivation of $H, FS \twoheadrightarrow H', FS'$ and part (1). Part (3) is proved by induction on the derivation of $H, P \rightsquigarrow H', P'$ and part (2).

□

References

[1] Gavin M. Bierman, Claudio V. Russo, Geoffrey Mainland, Erik Meijer, and Mads Torgersen. Pause 'n' play: Formalizing asynchronous C#. In *ECOOP*, volume 7313, pages 233–257. Springer, 2012.

[2] Patrick Th. Eugster, Pascal Felber, Rachid Guerraoui, and Anne-Marie Kermarrec. The many faces of publish/subscribe. *ACM Comput. Surv*, 35(2):114–131, 2003.

[3] Philipp Haller, Aleksandar Prokopec, Heather Miller, Viktor Klang, Roland Kuhn, and Vojin Jovanovic. Scala Improvement Proposal 14 - Futures and Promises. `http://docs.scala-lang.org/sips/pending/futures-promises.html`, 2012.

[4] Philipp Haller and Jason Zaugg. SIP-22: Async. `http://docs.scala-lang.org/sips/pending/async.html`, 2013.

$$\frac{}{\Gamma; \Delta \vdash \underline{b} : \texttt{Boolean}} \quad \text{(T-Boolean)}$$

$$\frac{}{\Gamma; \Delta \vdash \underline{i} : \texttt{Int}} \quad \text{(T-Int)}$$

$$\frac{x : \rho \in \Gamma}{\Gamma; \Delta \vdash x : \rho} \quad \text{(T-Var)}$$

$$\frac{}{\Gamma; \Delta \vdash \texttt{null} : \rho} \quad \text{(T-Null)}$$

$$\frac{\Gamma; \Delta \vdash x : \texttt{Boolean} \qquad \Gamma; \Delta \vdash e : \rho \qquad \Gamma; \Delta \vdash e' : \rho}{\Gamma; \Delta \vdash \texttt{if}\ (x)\ \{e\}\ \texttt{else}\ \{e'\} : \rho} \quad \text{(T-Cond)}$$

$$\frac{ftype(\sigma, f) = \tau}{\Gamma, x : \sigma; \Delta \vdash x.f : \tau} \quad \text{(T-Field)}$$

$$\frac{fields(C) = \bar{f} \qquad ftype(C, \bar{f}) = \bar{\tau} \qquad \Gamma; \Delta \vdash \bar{y} : \bar{\tau}}{\Gamma; \Delta \vdash \texttt{new}\ C(\bar{y}) : C} \quad \text{(T-New)}$$

$$\frac{mtype(\sigma, m) = (\bar{\tau}) \to \sigma' \qquad \Gamma, x : \sigma \vdash \bar{y} : \bar{\tau}}{\Gamma, x : \sigma; \Delta \vdash x.m(\bar{y}) : \sigma'} \quad \text{(T-Invoke)}$$

$$\frac{\Gamma; \Delta \vdash \bar{y} : \texttt{Observable[}\bar{\tau}\texttt{]} \qquad \Gamma; \sigma \circ \Delta \vdash e : \rho}{\Gamma; \Delta \vdash \texttt{rasync[}\sigma\texttt{]}(\bar{y})\ \{e\} : \texttt{Observable[}\sigma\texttt{]}} \quad \text{(T-RAsync)}$$

$$\frac{\Gamma; \Delta \vdash x : \texttt{Observable[}\sigma\texttt{]}}{\Gamma; \Delta \vdash \texttt{await}(x) : \sigma} \quad \text{(T-Await)}$$

$$\frac{\Gamma; \Delta \vdash e : \rho \qquad \Gamma, x : \rho; \Delta \vdash t : \sigma}{\Gamma; \Delta \vdash \texttt{let}\ x = e\ \texttt{in}\ t : \sigma} \quad \text{(T-Let)}$$

$$\frac{\Gamma; \Delta \vdash x : \texttt{Boolean} \qquad \Gamma; \Delta \vdash t : \rho}{\Gamma; \Delta \vdash \texttt{while}\ (x)\ \{t\} : \texttt{Boolean}} \quad \text{(T-While)}$$

$$\frac{\Gamma; \sigma \circ \Delta \vdash x : \sigma}{\Gamma; \sigma \circ \Delta \vdash \texttt{yield}(x) : \sigma} \quad \text{(T-Yield)}$$

Figure 11: Type checking expressions.

[5] Anders Hejlsberg, Mads Togersen, Scott Wiltamuth, and Peter Golde, editors. *The C# Programming Language*. Addison-Wesley, fourth edition, 2011.

[6] Erik Meijer. Your mouse is a database. *Commun. ACM*, 55(5):66–73, 2012.

[7] Don Syme, Tomas Petricek, and Dmitry Lomov. The F# asynchronous programming model. In Ricardo Rocha and John Launchbury, editors, *PADL*, pages 175–189. Springer, 2011.